

Bottom- k and Priority Sampling, Set Similarity and Subset Sums with Minimal Independence*

Mikkel Thorup

AT&T Labs—Research and University of Copenhagen
mikkel2thorup@gmail.com

March 25, 2013

Abstract

We consider bottom- k sampling for a set X , picking a sample $S_k(X)$ consisting of the k elements that are smallest according to a given hash function h . With this sample we can estimate the frequency $f = |Y|/|X|$ of any subset Y as $|S_k(X) \cap Y|/k$. A standard application is the estimation of the Jaccard similarity $f = |A \cap B|/|A \cup B|$ between sets A and B . Given the bottom- k samples from A and B , we construct the bottom- k sample of their union as $S_k(A \cup B) = S_k(S_k(A) \cup S_k(B))$, and then the similarity is estimated as $|S_k(A \cup B) \cap S_k(A) \cap S_k(B)|/k$.

We show here that even if the hash function is only 2-independent, the expected relative error is $O(1/\sqrt{fk})$. For $fk = \Omega(1)$ this is within a constant factor of the expected relative error with truly random hashing.

For comparison, consider the classic approach of repeated min-wise hashing, where we use k independent hash functions h_1, \dots, h_k , storing the smallest element with each hash function. For min-wise hashing, there can be a constant bias with constant independence, and this is not reduced with more repetitions k . Recently Feigenblat et al. showed that bottom- k circumvents the bias if the hash function is 8-independent and k is sufficiently large. We get down to 2-independence for any k . Our result is based on a simply union-bound, transferring generic concentration bounds for the hashing scheme to the bottom- k sample, e.g., getting stronger probability error bounds with higher independence.

For weighted sets, we consider priority sampling which adapts efficiently to the concrete input weights, e.g., benefiting strongly from heavy-tailed input. This time, the analysis is much more involved, but again we show that generic concentration bounds can be applied.

1 Introduction

The concept of min-wise hashing (or the “MinHash algorithm” according to ¹) is a basic algorithmic tool suggested by Broder et al. [6, 8] for problems related to set similarity and containment. After the initial application of this algorithm in the early Altavista search engine to detecting and clustering similar documents, the scheme has reappeared in numerous other applications¹ and is now a standard tool in data mining where it is used for estimating similarity [6, 8, 9], rarity [13], document duplicate detection [7, 20, 22, 37], etc [2, 4, 10, 30].

*A short version of this paper will appear at STOC’13.

¹See <http://en.wikipedia.org/wiki/MinHash>

In an abstract mathematical view, we have two sets, A and B , and we are interested in understanding their overlap in the sense of the Jaccard similarity $f = \frac{|A \cap B|}{|A \cup B|}$. In order to do this by sampling, we need sampling correlated between the two sets, so we sample by hashing. Consider a hash function $h : A \cup B \rightarrow [0, 1]$. For now we assume that h is fully random, and has enough precision that collisions are avoided with probability 1. The main mathematical observation is that $\Pr[\argmin h(A) = \argmin h(B)]$ is precisely $f = |A \cap B| / |A \cup B|$. Thus, we may sample the element with the minimal hash from each set, and use them in $[\argmin h(A) = \argmin h(B)]$ for an unbiased estimate of f . Here, for a logical statement S , $[S] = 1$ if S is true; otherwise $[S] = 0$.

For more concentrated estimators, we use repetition with k independent hash functions, h_1, \dots, h_k . For each set A , we store $M^k(A) = (\argmin h_1(A), \dots, \argmin h_k(A))$, which is a sample with replacement from A . The Jaccard similarity between sets A and B is now estimated as $|M^k(A) \cap M^k(B)|/k$ where $|M^k(A) \cap M^k(B)|$ denotes the number of agreeing coordinates between $M^k(A)$ and $M^k(B)$. We shall refer to this approach as repeated min-wise or $k \times \min$.

For our discussion, we consider the very related application where we wish to store a sample of a set X that we can use to estimate the frequency $f = \frac{|Y|}{|X|}$ of any subset $Y \subseteq X$. The idea is that the subset Y is not known when the sample from X is made. The subset Y is revealed later in the form of a characteristic function that can tell if (sampled) elements belong to Y . Using the $k \times \min$ sample $M^k(X)$, we estimate the frequency as $|M^k(X) \cap Y|/k$ where $|M^k(X) \cap Y|$ denotes the number of samples from $M^k(X)$ in Y .

Another classic approach for frequency estimation is to use just one hash function h and use the k elements from X with the smallest hashes as a sample $S_k(X)$. This is a sample without replacement from X . As in [12], we refer to this as a bottom- k sample. The method goes back at least to [19]. The frequency of Y in X is estimated as $|Y \cap S_k(X)|/k$. Even though surprisingly fast methods have been proposed to compute $k \times \min$ [3], the bottom- k signature is much simpler and faster to compute. In a single pass through a set, we only apply a single hash function h to each element, and use a max-priority queue to maintain the k smallest elements with respect to h .

It is standard¹ to use bottom- k samples to estimate the Jaccard similarity between sets A and B , for this is exactly the frequency of the intersection in the union. First we construct the bottom- k sample $S_k(A \cup B) = S_k(S_k(A) \cup S_k(B))$ of the union by picking the k elements from $S_k(A) \cup S_k(B)$ with the smallest hashes. Next we return $|S_k(A) \cap S_k(B) \cap S_k(A \cup B)|/k$.

Stepping back, for subset frequency, we generally assume that we can identify samples from the subset. In the application to set similarity, it is important that the samples are coordinated via hash functions, for this is what allows us to identify samples from the intersection as being sampled in both sets. In our mathematical analysis we will focus on the simpler case of subset frequency estimation, but in the application to set similarity that motivates our special interest in sampling via hash functions.

Limited independence The two approaches $k \times \min$ and bottom- k are similar in spirit, starting from the same base $1 \times \min = \text{bottom-1}$. With truly random hash functions, they have essentially the same relative standard deviation (standard deviation divided by expectation) bounded by $1/\sqrt{fk}$ where f is the set similarity or subset frequency. The two approaches are, however, very different from the perspective of pseudo-random hash functions of limited independence [36]: a random hash function h is d -independent if the hash values of any d given elements are totally random.

With min-wise hashing, we have a problem with bias in the sense of sets in which some elements have a better than average chance of getting the smallest hash value. It is known that $1 + o(1)$ bias requires $\omega(1)$ -independence [27]. This bias is not reduced by repetitions as in $k \times \min$. However, recently Porat et al. [18] proved that the bias for bottom- k vanishes for large enough $k \gg 1$ if we use 8-independent hashing.

Essentially they get an expected relative error of $O(1/\sqrt{fk})$, and error includes bias. For $fk = \Omega(1)$, this is only a constant factor worse than with truly random hashing. Their results are cast in a new framework of “ d - k -min-wise hashing”, and the translation to our context is not immediate.

Results In this paper, we prove that bottom- k sampling preserves the expected relative error of $O(1/\sqrt{fk})$ with 2-independent hashing, and this holds for any k including $k = 1$. We note that when $fk = o(1)$, then $1/\sqrt{fk} = \omega(1)$, so our result does not contradict a possible large bias for $k = 1$.

We remark that we also get an $O(1/\sqrt{(1-f)k})$ bound on the expected relative error. This is important if we estimate the dissimilarity $1 - f$ of sets with large similarity $f = 1 - o(1)$.

For the more general case of weighted sets, we consider priority sampling [17] which adapts near-optimally to the concrete input weights [32], e.g., benefiting strongly from heavy-tailed input. We show that 2-independent hashing suffices for good confidence intervals.

Our positive finding with 2-independence contrasts recent negative results on the insufficiency of low independence, e.g., that linear probing needs the 5-independence [27] that was proved sufficient by Pagh et al. [26].

Implementation For 2-independent hashing we can use the fast multiplication-shift scheme from [14], e.g., if the elements are 32-bit keys, we pick two random 64-bit numbers a and b . The hash of key x is computed with the C-code $(a * x + b) >> 32$, where $*$ is 64-bit multiplication which discards overflow, and $>>$ is a right shift. This is 10-20 times faster than the fastest known 8-independent hashing based on a degree 7 polynomial tuned for a Mersenne prime field [35]².

Practical relevance We note that Mitzenmacher and Vadhan [23] have proved that 2-independence generally works if the input has enough entropy. However, the real world has lots of low entropy data. In [35] it was noted how consecutive numbers with zero entropy made linear probing with 2-independent hashing extremely unreliable. This was a problem in connection with denial-of-service attacks using consecutive IP-addresses. For our set similarity, we would have similar issues in scenarios where small numbers are more common, hence where set intersections are likely to be fairly dense intervals of small numbers whereas the difference is more likely to consist of large random outliers. Figure 1 presents an experiment showing what happens if we try to estimate the dissimilarity with 2-independent hashing.

Stepping back, the result Mitzenmacher and Vadhan is that 2-independence works for sufficiently random input. In particular, we do not expect problems to show up in random tests. However, this does not imply that 2-independent hashing can be trusted on real data unless we have special reasons to believe that the input has high entropy. In Figure 1, bottom- k performs beautifully with 2-independent hashing, but no amount of experiments can demonstrate general reliability. However, the mathematical result of this paper is that bottom- k can indeed be trusted with 2-independent hashing: the expected relative error is $O(1/\sqrt{fk})$ no matter the structure of the input.

Techniques To appreciate our analysis, let us first consider the trivial case where we are given a non-random threshold probability p and sample all elements that hash below p . As in [16] we refer to this as *threshold sampling*. Since the hash of an element x is uniform in $[0, 1]$, this samples x with probability p . The sampling of x depends only on the hash value of x , so if, say, the hash function is d -independent, then the

²See Table 2 in [35] for comparisons with different key lengths and computers between multiplication-shift (TwoIndep), and tuned polynomial hashing (CWtrick). The table considers polynomials of degree 3 and 4, but the cost is linear in the degree, so the cost for degree 7 is easily extrapolated.

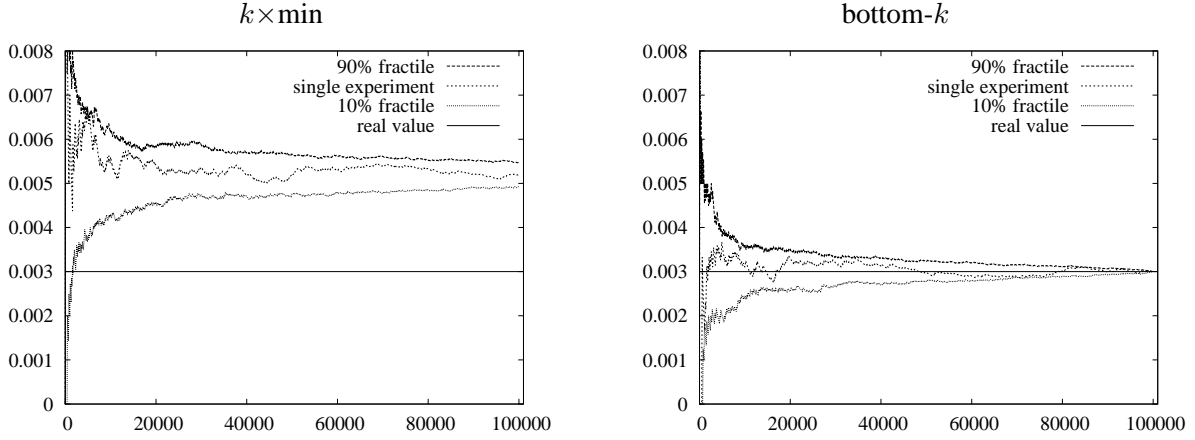


Figure 1: Experiment with set consisting of 100300 32-bit keys. It has a “core” consisting of the consecutive numbers 1, ...100000. In addition it has 300 random “outliers”. Using k samples from the whole set, we want to estimate the frequency of the outliers. The true frequency is $\frac{300}{100300} \approx 0.003$. We used $k = 1, \dots, 100000$ in $k \times \min$ and $\text{bottom-}k$ and made one hundred experiments. For each k , we sorted the estimates, plotting the 10th and 90th value, labeled as 10% and 90% fractile in the figures. We also plotted the results from a single experiment. For readability, only one in every 100 values of k is plotted. Both schemes converge, but due to bias, $k \times \min$ converges to a value that is 70% too large. Since $\text{bottom-}k$ does sampling without replacement, it becomes exact when the number of samples is the size of the whole set. The bias is a function of the structure of the subset within the whole set, e.g., the core set must have a negative bias complementing the positive bias of the outliers. It is therefore not possible to correct for the bias if one only has the sample available.

number of samples is the sum of d -independent 0-1 variables. This scenario is very well understood (see, e.g., [15, 31]).

We could set $p = k/n$, and get an expected number of k samples. Morally, this should be similar to a bottom- k sample, which is what we get if we end up with exactly k samples, that is, if we end up with $h_{(k)} < p \leq h_{(k+1)}$ where $h_{(i)}$ denotes the i th smallest hash value. What complicates the situation is that $h_{(k)}$ and $h_{(k+1)}$ are random variables depending on all the random hash values.

An issue with threshold sampling is that the number of samples is variable. This is an issue if we have bounded capacity to store the samples. With k expected samples, we could put some limit $K \gg k$ on the number of samples, but any such limit introduces dependencies that have to be understood. Also, if we have room for K samples, then it would seem wasteful not fill it with a full bottom- K sample.

Our analysis of bottom- k samples is much simpler than the one in [18] for 8-independent hashing with $k \gg 1$. With a union bound we reduce the analysis of bottom- k samples to the trivial case of threshold sampling. Essentially we only get a constant loss in the error probabilities. With 2-independent hashing, we then apply Chebyshev's inequality to show that the expected relative error is $O(1/\sqrt{fk})$. The error probability bounds are immediately improved if we use hash functions with higher independence.

It is already known from [5] that we can use a 2-independent bottom- k sample of a set to estimate its size n with an expected error of $O(\sqrt{n})$. The estimate is simply the inverse of the k th smallest sample. Applying this to two $\Theta(n)$ -sized sets and their union, we can estimate $|A|$, $|B|$, $|A \cup B|$ and $|A \cap B| = |A| + |B| - |A \cup B|$ each with an expected error of $O(\sqrt{n})$. However, $|A \cap B|$ may be much smaller than $O(\sqrt{n})$. If we instead multiply our estimate of the similarity $f = |A \cap B|/|A \cup B|$ with the estimate of $|A \cup B|$, the resulting estimate of $|A \cap B|$ is

$$(1 \pm O(1/\sqrt{fk}))f(|A \cup B| \pm O(\sqrt{n})) = |A \cap B| \pm O(\sqrt{|A \cap B|}).$$

The analysis of priority sampling for weighted sets is much more delicate, but again, using union bounds, we show that generic concentration bounds apply.

2 Bottom- k samples

We are given a set of X of n elements. A hash function maps the elements uniformly and collision free into $[0, 1]$. Our bottom- k sample S consists of the k elements with the lowest hash values. The sample is used to estimate the frequency $f = |Y|/|X|$ of any subset Y of X as $|Y \cap S|/k$. With 2-independent hashing, we will prove the following error probability bound for any $r \leq \bar{r} = \sqrt{k}/3$:

$$\Pr \left[\left| |Y \cap S| - fk \right| > r\sqrt{fk} \right] \leq 4/r^2. \quad (1)$$

The result is obtained via a simple union bound where stronger hash functions yield better error probabilities. With d -independence with d an even constant, the probability bound is $O(1/r^d)$.

It is instructive to compare d -independence with the idea of storing d independent bottom- k samples, each based on 2-independence, and use the median estimate. Generally, if the probability of a certain deviation is p , the deviation probability for the median is bounded by $(2ep)^{d/2}$, so the $4/r^2$ from (1) becomes $(2e4/r^2)^{d/2} < (5/r)^d$, which is the same type of probability that we get with a single d -independent hash function. The big advantage of a single d -independent hash function is that we only have to store a single bottom- k sample.

If we are willing to use much more space for the hash function, then we can use twisted tabulation hashing [28] which is very fast, and then we get exponential decay in r though only down to an arbitrary polynomial the space used.

In order to show that the expected relative error is $O(1/\sqrt{fk})$, we also prove the following bound for $fk \leq 1/4$:

$$\Pr[|Y \cap S| \geq \ell] = O(fk/\ell^2 + \sqrt{f}/\ell). \quad (2)$$

From (1) and (2) we get

Proposition 1 *For bottom- k samples based on 2-independent hashing, a fraction f subset is estimated with an expected relative error of $O(1/\sqrt{fk})$.*

Proof The proof assumes (1) and (2). For the case $fk > 1/4$, we will apply (1). The statement is equivalent to saying that the sample error $||Y \cap S| - fk|$ in expectation is bounded by $O(\sqrt{fk})$. This follows immediately from (1) for errors below $\bar{r}\sqrt{fk} = k\sqrt{|Y|/n}$. However, by (1), the probability of a larger error is bounded by $4/\bar{r}^2 = O(1/k)$. The maximal error is k , so the contribution to the expected error is $O(1)$. This is $O(\sqrt{fk})$ since $fk > 1/4$.

We will now handle the case $fk \leq 1/4$ using (2). We want to show that the expected absolute error is $O(\sqrt{fk})$. We note that only positive errors can be bigger than fk , so if the expected error is above $2fk$, the expected number of samples from Y is proportional to the expected error. We have $\sqrt{fk} \geq 2fk$, so for the expected error bound, it suffices to prove that the expected number of samples is $|Y \cap S| = O(\sqrt{fk})$. Using (2) for the probabilities, we now sum the contributions over exponentially increasing sample sizes.

$$\begin{aligned} \mathbb{E}[|Y \cap S|] &\leq \sum_{i=0}^{\lfloor \lg k \rfloor} (2^{i+1} \Pr[|Y \cap S| \geq 2^i]) \\ &= \sum_{i=0}^{\lfloor \lg k \rfloor} O\left(2^i(fk/2^{2i} + \sqrt{f}/2^i)\right) \\ &= O\left(fk + \sqrt{f}(1 + \lg k)\right) = O\left(\sqrt{fk}\right). \end{aligned}$$

■

2.1 A union upper bound

First we consider overestimates. For positive parameters a and b to be chosen, we will bound the probability of the overestimate

$$|Y \cap S| > \frac{1+b}{1-a} fk. \quad (3)$$

Define the threshold probability

$$p = \frac{k}{n(1-a)}.$$

Note that p is defined deterministically, independent of any samples. It is easy to see that the overestimate (3) implies one of the following two threshold sampling events:

- (A) The number of elements from X that hash below p is less than k . We expected $pn = k/(1-a)$ elements, so k is a factor $(1-a)$ below the expectation.
- (B) Y gets more than $(1+b)p|Y|$ hashes below p , that is, a factor $(1+b)$ above the expectation.

To see this, assume that both (A) and (B) are false. When (A) is false, we have k hashes from X below p , so the largest hash in S is below p . Now if (B) is also false, we have at most $(1+b)p|Y| = (1+b)/(1-a) \cdot fk$ elements from Y hashing below p , and only these elements from Y could be in S . This contradicts (3). By the union bound, we have proved

Proposition 2 *The probability of the overestimate (3) is bounded by $P_A + P_B$ where P_A and P_B are the probabilities of the events (A) and (B), respectively.*

Upper bound with 2-independence Addressing events like (A) and (B), let m be the number of elements in the set Z considered, e.g., $Z = X$ or $Z = Y$. We study the number of elements hashing below a given threshold $p \in [0, 1]$. Assuming that the hash values are uniform in $[0, 1]$, the mean is $\mu = mp$. Assuming 2-independence of the hash values, the variance is $mp(1-p) = (1-p)\mu$ and the standard deviation is $\sigma = \sqrt{(1-p)\mu}$. By Chebyshev's inequality, we know that the probability of a deviation by $r\sigma$ is bounded by $1/r^2$. Below we will only use that the relative standard deviation σ bounded by $1/\sqrt{\mu}$.

For any given $r \leq \sqrt{k}/3$, we will fix a and b to give a combined error probability of $2/r^2$. More precisely, we will fix $a = r/\sqrt{k}$ and $b = r/\sqrt{fk}$. This also fixes $p = k/(n(1-a))$. We note for later that $a \leq 1/3$ and $a \leq b$. This implies

$$(1+b)/(1-a) \leq (1+3b) = 1 + 3r/\sqrt{fk}. \quad (4)$$

In connection with (A) we study the number of elements from X hashing below p . The mean is $pn \geq k$ so the relative standard deviation is less than $1/\sqrt{k}$. It follows that a relative error of $a = r/\sqrt{k}$ corresponds to at least r standard deviations, so

$$P_A = \Pr[\#\{x \in X | h(x) < p\} < (1-a)np] < 1/r^2.$$

In connection with (B) we study the number of elements from Y hashing below p . Let $m = |Y|$. The mean is $pm = km/(n(1-a))$ and the relative standard deviation less than $1/\sqrt{pm} < 1/\sqrt{km/n}$. It follows that a relative error of $b = r/\sqrt{km/n}$ is more than r standard deviations, so

$$P_B = \Pr[\#\{y \in Y | h(y) < p\} > (1+b)mp] < 1/r^2.$$

By Proposition 2 we conclude that the probability of (3) is bounded by $2/r^2$. Rewriting (3) with (4), we conclude that

$$\Pr[|Y \cap S| > fk + 3r\sqrt{fk}] \leq 2/r^2. \quad (5)$$

This bounds the probability of the positive error in (1). The above constants 3 and 2 are moderate, and they can easily be improved if we look at asymptotics. Suppose we want good estimates for subsets Y of frequency at least f_{\min} , that is, $|Y| \geq f_{\min}|X|$. This time, we set $a = r/\sqrt{fk}$, and then we get $P_A \leq f/r^2$. We also set $b = r/\sqrt{fk}$ preserving $P_B \leq 1/r^2$. Now for any $Y \subseteq X$ with $|Y| > fn$, we have

$$\Pr[|Y \cap S| > (1+\varepsilon)fk] = (1+f)/r^2 \quad (6)$$

$$\text{where } \varepsilon = \frac{1+r/\sqrt{fk}}{1-r/\sqrt{k}} - 1 = \frac{r/\sqrt{k} + r/\sqrt{fk}}{1-r/\sqrt{k}}.$$

With $f = o(1)$ and $k = \omega(1)$, the error is $\varepsilon = (1+o(1))r/\sqrt{fk}$, and the error probability is $P_\varepsilon = (1+f)/r^2 = (1+o(1))/r^2$. Conversely, this means that if we for subsets of frequency f and a relative positive error ε want an error probability around P_ε , then we set $r = \sqrt{1/P_\varepsilon}$ and $k = r^2/(f\varepsilon^2) = 1/(fP_\varepsilon\varepsilon^2)$.

2.2 A union lower bound

We have symmetric bounds for underestimates:

$$|Y \cap S| < \frac{1 - b'}{1 + a'} f k. \quad (7)$$

This time we define the threshold probability $p' = \frac{k}{n(1+a')}$. It is easy to see that the overestimate (3) implies one of the following two events:

(A') The number of elements from X below p' is at least k . We expected $p'n = k/(1 + a')$ elements, so k is a factor $(1 + a')$ above the expectation.

(B') Y gets less than $(1 - b')p|Y|$ hashes below p' , that is, a factor $(1 - b')$ below the expectation.

To see this, assume that both (A') and (B') are false. When (A') is false, we have less than k hashes from X below p' , so S must contain all hashes below p' . Now if (B) is also false, we have at least $(1 - b)p'|Y| = (1 - b)/(1 + a) \cdot f k$ elements from $Y \subseteq X$ hashing below p' , hence which must be in S . This contradicts (7). By the union bound, we have proved

Proposition 3 *The probability of the underestimate (7) is bounded by $P_{A'} + P_{B'}$ where $P_{A'}$ and $P_{B'}$ are the probabilities of the events (A') and (B'), respectively.*

Lower bound with 2-independence Using Proposition 3 we will bound the probability of underestimates, complementing our previous probability bounds for overestimates from Section 2.1. We will provide bounds for the same overall relative error as we did for the overestimates; namely

$$\varepsilon = \frac{1 + b}{1 - a} - 1 = (a + b)/(1 - a)$$

However, for the events (A') and (B') we are going to scale up the relative errors by a factor $(1 + a)$, that is, we will use $a' = a(1 + a)$ and $b' = b(1 + a)$. The overall relative negative error from (7) is then

$$\begin{aligned} \varepsilon' &= 1 - \frac{1 - b'}{1 + a'} = (a' + b')/(1 + a') \\ &< (1 + a)(a + b)/(1 + a') < (a + b) < \varepsilon. \end{aligned}$$

Even with this smaller error, we will get better probability bounds than those we obtained for the overestimates. For (A) we used $1/\sqrt{k}$ as an upper bound on the relative standard deviation, so a relative error of a was counted as $s_A = a\sqrt{k}$ standard deviations. In (A') we have mean $\mu' = np' = k/(1 + a')$, so the relative standard deviation is bounded by $1/\sqrt{k/(1 + a')} = \sqrt{1 + a + a^2}/\sqrt{k}$. This means that for (A'), we can count a relative error of $a' = a(1 + a)$ as

$$\begin{aligned} s'_A &= a(1 + a)\sqrt{k}/\sqrt{1 + a + a^2} \\ &= s_A(1 + a)/\sqrt{1 + a + a^2} > s_A \end{aligned}$$

standard deviations. In Section 2.1 we bounded P_A by $1/s_A^2$, and now we can bound $P_{A'}$ by $1/s'^2_A \leq 1/s_A^2$. The scaling has the same positive effect on our probability bounds for (B'). That is, in Section 2.1, a relative

error of b was counted as $s_B = b\sqrt{fk}$ standard deviations. With (B') our relative error of $b' = b(1+a)$ is counted as

$$\begin{aligned} s'_B &= b(1+a)\sqrt{fk}/\sqrt{1+a+a^2} \\ &= s_B(1+a)/\sqrt{1+a+a^2} > s_B \end{aligned}$$

standard deviations, and then we can bound $P_{B'}$ by $1/s_B'^2 \leq 1/s_B^2$. Summing up, our negative relative error ε' is smaller than our previous positive error ε , and our overall negative error probability bound $1/s_A'^2 + 1/s_B'^2$ is smaller than our previous positive error probability bound $1/s_A^2 + 1/s_B^2$. We therefore translate (5) to

$$\Pr[|Y \cap S| < fk - 3r\sqrt{fk}] \leq 2/r^2. \quad (8)$$

which together with (5) establishes (1). Likewise (6) translates to

$$\begin{aligned} \Pr[||Y \cap S| - fk| > \varepsilon fk] &\leq 2(1+f)/r^2 \\ \text{where } \varepsilon &= \frac{1+r/\sqrt{fk}}{1-r/\sqrt{k}} - 1. \end{aligned} \quad (9)$$

As for the positive error bounds we note that with $f = o(1)$ and $k = \omega(1)$, the error is $\varepsilon = (1+o(1))r/\sqrt{fk}$ and the error probability is $P_\varepsilon = (2+o(1))/r^2$. Conversely, this means that if we for a target relative error ε want an error probability around P_ε , then we set $r = \sqrt{2/P_\varepsilon}$ and $k = r^2/(f\varepsilon^2) = 2/(fP_\varepsilon\varepsilon^2)$.

2.3 Rare subsets

We now consider the case where the expected number fk of samples from Y is less than $1/4$. We wish to prove (2)

$$\Pr[|Y \cap S| \geq \ell] = O(fk/\ell^2 + \sqrt{f}/\ell).$$

For some balancing parameter $c \geq 2$, we use the threshold probability $p = ck/n$. The error event (A) is that less than k elements from X sample below p . The error event (B) is that at least ℓ elements hash below p . As in Proposition 2, we observe that ℓ bottom- k samples from Y implies (A) or (B), hence that $\Pr[|Y \cap S| \geq \ell] \leq P_A + P_B$.

The expected number of elements from X that hash below p is ck . The error event (A) is that we get less than k , which is less than half the expectation. This amounts to at least $\sqrt{ck}/2$ standard deviations, so by Chebyshev's inequality, the probability of (A) is $P_A \leq 1/(\sqrt{ck}/2)^2 = 4/(ck)$.

The event (B) is that at least ℓ elements from Y hash below p , while the expectation is only fck . Assuming that $\ell \geq 2fck$, the error is by at least $(\ell/2)/\sqrt{fck}$ standard deviations. By Chebyshev's inequality, the probability of (B) is $P_B \leq 1/((\ell/2)/\sqrt{fck})^2 = 4fck/\ell^2$. Thus

$$P_A + P_B \leq 4/(ck) + 4fck/\ell^2.$$

We wish to pick c for balance, that is,

$$4/ck = 4fck/\ell^2 \iff c = \ell/(\sqrt{fk})$$

However, we have assumed that $c \geq 2$ and that $\ell \geq 2fck$. The latter is satisfied because $2fck = 2fk\ell/(\sqrt{fk}) = 2\sqrt{f}\ell$ and $f \leq 1/4$. Assuming that $c = \ell/(\sqrt{fk}) \geq 2$, we get

$$P_A + P_B \leq 8/(k(\ell/(\sqrt{fk}))) = 8\sqrt{f}/\ell.$$

When $\ell/(\sqrt{f}k) < 2$, we set $c = 2$. Then

$$P_A + P_B \leq 2/k + 8fk/\ell^2 \leq 16fk/\ell^2.$$

Again we need to verify that $\ell \geq fck = 2fk$, but that follows because $\ell \geq 1$ and $fk \leq 1/4$. We know that at one of the above two cases applies, so we conclude that

$$P[|Y \cap S| \geq \ell] \leq P_A + P_B = O(fk/\ell^2 + \sqrt{f}/\ell),$$

completing the proof of (2).

3 Priority sampling

We now consider the more general situation where we are dealing with a set I of weighted items with w_i denoting the weight of item $i \in I$. Let $\sum I = \sum_{i \in I} w_i$ denote the total weight of set I .

Now that we are dealing with weighted items, we will use *priority sampling* [17] which generalizes the bottom- k samples we used for unweighted elements. The input is a set of I of weighted items. Each item or element i is identified by a unique key which is hashed uniformly to a random number $h_i \in (0, 1)$. The item is assigned a *priority* $q_i = w_i/h_i > w_i$. We assume that all priorities end up distinct and different from the weights. If not, we could break ties based on an ordering of the items. The priority sample S of size k contains the k samples of highest priority, but it also stores a threshold τ which is the $(k + 1)$ th highest priority. Based on this we assign a weight estimate \hat{w}_i to each item i . If i is not sampled, $\hat{w}_i = 0$; otherwise $\hat{w}_i = \max\{w_i, \tau\}$. A basic result from [17] is that $E[\hat{w}_i] = w_i$ if the hash function is truly random (in [17], the h_i were described as random numbers, but here they are hashes of the keys).

We note that priority sampling generalize the bottom- k sample we used for unweighted items, for if all weights are unit, then the k highest priorities correspond to the k smallest hash values. In fact, priority sampling predates [12], and [12] describes bottom- k samples for weighted items as a generalization of priority sampling, picking the first k items according to an arbitrary randomized function of the weights.

The original objective of priority sampling [17] was subset sum estimation. A subset $J \subseteq I$ of the items is selected, and we estimate the total weight in the subset as $\hat{w}_J = \sum\{\hat{w}_i | i \in J \cap S\}$. By linearity of expectation, this is an unbiased estimator. A cool application from [17] was that as soon as the signature of the Slammer worm [24] was identified, we could inspect the priority samples from the past to track its history and identify infected hosts. An important point is that the Slammer worm was not known when the samples were made. Samples are made with no knowledge on which subsets will later turn out to be of interest.

Trivially, if we want to estimate the relative subset weight $\sum J / \sum I$ and we do not know the exact total, we can divide \hat{w}_J with the estimate \hat{w}_I of the total. As with the bottom- k sampling for unweighted items, we can easily use priority sampling to estimate the similarity of sets of weighted items: given the priority sample from two sets, we can easily construct the priority sample of their union, and estimate the intersection as a subset. This is where it is important that we use a hash function so that the sampling from different sets is coordinated, e.g., we could not use iterative sampling procedures like the one in [11]. In the case of histogram similarity, it is natural to allow the same item to have different weights in different sets. More specifically, allowing zero weights, every possible item has a weight in each set. For the similarity we take the sum of the minimum weight for each item, and divide it by the sum of the maximum weight for each item. This requires a special sampling that we shall return to at the end.

Priority sampling is not only extremely easy to implement on-line with a standard min-priority queue; it also has some powerful universal properties in its adaption to the concrete input weights. As proved in [32],

given one extra sample, priority sampling has smaller variance sum $\sum_i \text{Var}[\hat{w}_i]$ than *any* off-line sampling scheme tailored for the concrete input weights. In particular, priority sampling benefits strongly if there are dominant weights w_i in the input, estimated precisely as $\hat{w}_i = \max\{w_i, \tau\} = w_i$. In the important case of heavy tailed input distributions [1], we thus expect most of the total weight to be estimated without any error. The quality of a priority sample is therefore often much better than what can be described in terms of simple parameters such as total weight, number of items, etc. The experiments in [17] on real and synthetic data show how priority sampling sometimes gains orders of magnitude in estimate quality over competing methods.

The quality of a priority estimate depends completely on the distribution of weights in the input, and often we would like to know how much we can trust a given estimate. What we really want from a sample is not just an estimate, but a confidence interval for the subset sum [33]: from the sample we want to compute lower and upper bounds that capture the true value with some desired probability. The confidence interval does not have to be a simple nice function. It has to be efficiently computable, and we want it to be as tight as possible.

All the current analysis of priority sampling [17, 32, 33] is heavily based on true randomness, assuming that the priorities are independent random variables, e.g., the proof from [17] that $E[\hat{w}_i] = w_i$ starts by fixing the priorities q_j of all the other items $j \neq i$. However, in this paper, we want to use hash functions with independence as low as 2, and then any such approach breaks down. Estimates will no longer be unbiased, but we will show that good confidence intervals can still be computed.

To the best of our knowledge, our paper is the first to show that anything useful can be said about subset sum estimates based on given number of k samples made using < 8 -independent random variables. The essence of our result is that priority sampling gets reasonable performance with any hashing scheme as long as it is 2-independent.

Below we first develop error probability bounds for priority sampling with limited independence applied to a given set of input weights. Afterwards we show how confidence intervals can be derived from a sample. At the very end, we show how to estimate histogram similarity.

3.1 Threshold sampling

Generalizing the pattern for unweighted sets, our basic goal is to relate the error probabilities with priority sampling to the much simpler case of threshold sampling for weighted items [16]. In our description of threshold sampling, we will develop notation and basic results that will later be used for priority sampling.

In threshold sampling, we do not have a predefined sample size. Instead we have a given threshold t . We will still use exactly the same random priorities as above, but now an item is sampled if and only if $q_i > t$. The weight estimate is

$$\hat{w}_i^t = \begin{cases} 0 & \text{if } q_i \leq t \\ \max\{w_i, t\} & \text{if } q_i > t \end{cases} \quad (10)$$

Thus, if priority sampling leads to threshold τ , then the priority estimates are $\hat{w}_i = \hat{w}_i^\tau$. We shall use I^t to denote the items sampled with threshold t , that is, items i with $q_i > t$. The priority threshold τ is the $(k+1)$ st largest priority, hence the smallest value such that $|I^\tau| \leq k$.

We note that threshold sampling is well-known from statistics (see, e.g., [29]). With fully random and independent h_i , each item i is sampled independently: a so-called a Poisson sampling scheme. With the given weights w_i and threshold t , let k_t be the expected number of samples with threshold t . Among all possible Poisson sampling schemes with unbiased estimators \hat{w}_i and an expected number of k_t samples, threshold sampling is the unique scheme that minimizes the variance sum $\sum_{i \in I} \text{Var}[\hat{w}_i]$ [29, p 86].

Fractional subsets and inner products It is convenient to generalize from regular subsets to fractional subsets. For each $i \in I$, there is a fraction $f_i \in [0, 1]$. We want to estimate $f w$ denoting the inner product $\sum_{i \in I} f_i w_i$. We estimate it as $f \hat{w}^t$ denoting $\sum_{i \in I} f_i \hat{w}_i^t = \sum_{i \in S} f_i \hat{w}_i^t$. Note that for this estimate, we only need to know f_i for $i \in S$. To emulate a standard subset J , we let f be the characteristic function of J , that is, $f_i = 1$ if $i \in J$; otherwise $f_i = 0$ otherwise. Using inner products will simplify a lot of notation in our analysis. The generalization to fractional subsets comes for free in our analysis which is all based on concentration bounds for sums of random variables $X_i \in [0, 1]$.

Notation for smaller weights Whenever we with threshold or priority sampling end up with a threshold t , we know that variability in the estimate is from items i with weight below t . We will generally use a subscript $<_t$ to denote the restriction to items i with weights $w_i < t$, e.g., $I_{<_t} = \{i \in I \mid w_i < t\}$ and $w_{<_t} = (w_i)_{i \in I_{<_t}}$ is the vector of these smaller weights. Then $f w_{<_t} = \sum_{i \in I_{<_t}} f_i w_i$. Notice that $f w_{<_t}$ does not include i with $w_i \geq t$ even if $f_i w_i < t$.

Above we defined $w_{<_t}$ to denote the vector $(w_i)_{i \in I_{<_t}}$ of weights below t , and used it for the inner product $f w_{<_t} = \sum_{i \in I_{<_t}} f_i w_i$. When it is clear from the context that we need a number, not a vector, we will use $w_{<_t}$ to denote the sum of these weights, that is, $w_{<_t} = \sum_{i \in I_{<_t}} w_i = \mathbf{1} w_{<_t}$ where $\mathbf{1}$ is the all 1s vector. Since $f_i \leq 1$ for all i , we always have $f w_{<_t} \leq w_{<_t}$.

We shall use subscript \leq_t , \geq_t , and $>_t$ to denote the corresponding restriction to items with weight $\leq t$, $\geq t$, and $> t$, respectively.

For a subset J of I , we let $f J = \sum_{i \in J} f_i$, thus identifying J with its characteristic vector. As an example, we can write our estimate with threshold t as

$$f \hat{w}^t = f w_{\geq t} + t(f I_{<_t}^t). \quad (11)$$

Error probability functions As in the unweighted case, the point in relating to threshold sampling is that error probability bounds for threshold sampling are easily derived. For items with unit weights, we reduced the bottom- k sampling error event to the union of four threshold sampling error events (A), (B), (A'), and (B'). However, now with weighted items, we are going to reduce the priority sampling error event to the union of an unbounded number of threshold sampling error events that happen with geometrically decreasing probabilities. Our reduction will hold for most hash functions, including 2-independent hash functions, but to make such a claim clear, we have to carefully describe what properties of the hash functions we rely on.

Assume that the threshold t is fixed. With reference to (11), the variability in our estimate is all from $f I_{<_t}^t$. An item $i \in I_{<_t}$ is sampled and included in $I_{<_t}^t$ if $q_i = w_i/h_i > t \iff h_i < w_i/t$, hence with probability w_i/t . If i is sampled, it adds $f_i \in [0, 1]$ to $f I_{<_t}^t$; otherwise 0.

The above is an instance of bounded random variables $X_i \in [0, 1]$, $i \in I_{<_t}$, where each X_i is a function of h_i ; namely $X_i = f_i[h_i < w_i/t]$. With the fixed threshold t , X_i depends only on the hash h_i . Therefore, if the hash values h_i are, say d -independent, then so are the X_i . Let $X = \sum_{i \in I_{<_t}} X_i$ and $\mu = E[X]$. We are interested in an error probability function \wp such that for $\mu > 0$, $\delta > 0$, if $\mu = E[X]$, then

$$\Pr[|X - \mu| > \delta \mu] \leq \wp(\mu, \delta).$$

The error probability function \wp that we can use depends on the quality of the hash function. For example, if the hash function is 2-independent, then $\text{Var}[X] \leq \mu$, and then by Chebyshev's inequality, we can use

$$\wp^{\text{Chebyshev}}(\mu, \delta) = 1/(\delta^2 \mu). \quad (12)$$

In the case of full randomness, for $\delta \leq 1$, we could use a standard 2-sided Chernoff bound (see, e.g., [25])

$$\wp^{\text{Chernoff}_{\delta \leq 1}}(\mu, \delta) = 2 \exp(-\delta^2 \mu / 3). \quad (13)$$

For most of our results, it is more natural to think of δ as a function of μ and some target error probability $P \in (0, 1)$, defining $\delta(\mu, P)$ such that

$$\mu(\mu, \delta(\mu, P)) = P. \quad (14)$$

Returning to threshold sampling with threshold t , for $i \in I_{< t}$, we have $X_i = f_i[i \in I_{< t}^t]$, $X = f I_{< t}^t$, and $\mu = f w_{< t} / t$. Moreover, by (11), $f \hat{w}^t - f w = t(f I_{< t}^t) - f w_{< t}$. Hence

$$\Pr[|f \hat{w}^t - f w| > \delta(f w_{< t} / t, P) f w_{< t}] \leq P. \quad (15)$$

When we start analyzing priority sampling, we will need to relate the probabilities of different threshold sampling events. This places some constraints on the error probability function \wp . Mathematically, it is convenient to allow \wp to attain values above 1, but only values below 1 are probabilistically interesting.

Definition 4 An error probability function $\wp : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is well-behaved if

- (a) \wp is continuous and strictly decreasing in both arguments.
- (b) If with the same absolute error we decrease the expectancy, then the probability goes down. Formally if $\mu' < \mu$ and $\mu' \delta' \geq \mu \delta$, then $\wp(\mu', \delta') < \wp(\mu, \delta)$.

We also have an optional condition for cases where we only care for $\delta \leq 1$ as in (13)

- (c) If $\delta \leq 1$ and $\wp(\mu, \delta) < P_{\wp}^{(c)}$ for some constant P_{\wp} depending on w_p , then $\wp(\mu, \delta)$ falls at least inversely proportional to $\mu \delta^2$. Formally, if $\delta_0, \delta_1 \leq 1$, $\wp(\mu_0, \delta_0) < 1$, and $\mu_0 \delta_0^2 < \mu_1 \delta_1^2$, then

$$\wp(\mu_0, \delta_0) \geq \frac{\mu_0 \delta_0^2}{\mu_1 \delta_1^2} \wp(\mu_1, \delta_1). \quad (16)$$

We will use condition (c) to argue that probabilities of different events fall geometrically. The condition is trivially satisfied with our 2-independent Chebyshev bound (12), so we can just set $P_{\wp}^{(c)} = 1$. The restrictions in (c) are necessary for the Chernoff bound (13), $\wp^{\text{Chernoff}_{\delta \leq 1}} = 2 \exp(-\delta^2 \mu / 3)$, which only falls quickly enough for $\delta^2 \mu / 3 \geq 1$, hence with $\wp^{\text{Chernoff}_{\delta \leq 1}}(\mu, \delta) \leq P_{\wp}^{(c)} = 2/e$. As a further illustration, with 4-independence, we have the 4th moment bound $\frac{\mu + 3\mu^2}{(\delta\mu)^4}$ (see, e.g., [21, Lemma 4.19]). For $\mu \geq 1$, this is upper bounded by $\wp^{\text{4th-moment}_{\mu \geq 1}}(\mu, \delta) = \left(\frac{2}{\mu \delta^2}\right)^2$. For $\delta \leq 1$, the condition $\mu \geq 1$ is satisfied if $\wp^{\text{4th-moment}_{\mu \geq 1}}(\mu, \delta) \leq 4$, so we can just use $P_{\wp}^{(c)} = 1$.

As an application of (a) and (b) we get

Lemma 5 For thresholds t', t , and relative errors δ', δ , if $t' < t$ and $\wp(f w_{< t'} / t', \delta_{< t'}) = \wp(f w_{< t} / t, \delta_{< t})$, then

$$\delta' f w_{< t'} < \delta f w_{< t}.$$

Hence, for any fixed target error probability P in (15), the target error

$$\delta(f w_{< t} / t, P) f w_{< t}$$

is strictly decreasing in the threshold t .

Proof We will divide the decrease from t to t' into a series of atomic decreases. The first atomic “decrease” is from $fw_{\leq t}$ to $fw_{< t}$. This makes no difference unless there are weights equal to t so that $fw_{< t} < fw_{\leq t}$. Assume this is the case and suppose $\wp(fw_{< t}/t, \delta_{< t}) = \wp(fw_{\leq t}/t, \delta_{\leq t})$. Since $fw_{\leq t}/t < fw_{< t}/t$, it follows directly from (b) that $\delta_{< t}fw_{< t}/t < \delta_{\leq t}fw_{\leq t}/t$, hence that $\delta_{< t}fw_{< t} < \delta_{\leq t}fw_{\leq t}$.

The other atomic decrease we consider is from $fw_{< t}$ to $fw_{\leq t'}$ where $t' < t$ and with no weights in (t', t) , hence with $fw_{\leq t'} = fw_{< t}$. Suppose $\wp(fw_{\leq t'}/t', \delta_{\leq t'}) = \wp(fw_{< t}/t, \delta_{< t})$. Since $t' < t$, $fw_{\leq t'}/t' > fw_{< t}/t$, so by (a), $\delta_{\leq t'} < \delta_{< t}$. It follows that $\delta_{\leq t'}fw_{\leq t'} < \delta_{< t}fw_{< t}$. Alternating between these two atomic decreases, we can implement an arbitrary decrease in the threshold as required for the lemma. ■

Threshold confidence intervals In the case of threshold sampling with a fixed threshold t , it is now trivial to derive confidence intervals for the true value fw . The sample gives us the exact value $fw_{\geq t}$ for weights at least as big as t , and an estimate $f\hat{w}_{< t}^t$ for those below. Setting

$$\begin{aligned} f\hat{w}_{< t}^- &= \min\{x \mid (1 + \delta(x/t, P))x \geq f\hat{w}_{\geq t}^t\} \\ f\hat{w}_{< t}^+ &= \max\{x \mid (1 - \delta(x/t, P))x \leq f\hat{w}_{\geq t}^t\} \end{aligned}$$

we get

$$\Pr[fw_{\geq t} + f\hat{w}_{< t}^- \leq fw \leq fw_{\geq t} + f\hat{w}_{< t}^+] \geq 1 - 2P.$$

3.2 Priority sampling: the main result

We are now ready to present our main technical result.

Theorem 6 For items $i \in I$, let be given a weight vector $(w_i)_{i \in I}$ with corresponding fractions $(f_i)_{i \in I}$. Let the error probability function \wp satisfy Definition 4 including (c). With a given target error probability $P \leq P_{\wp}^{(c)}$ and sample size k , consider a random priority sampling event, assigning to each item $i \in I$, a hash $h_i \in (0, 1)$ and priority $q_i = w_i/h_i$. Let τ be the resulting priority threshold, i.e, the $k + 1$ th largest priority. Let

$$\delta = 6\delta(f\hat{w}_{< \tau}/(3\tau), P).$$

If $\delta < 2$, then

$$\Pr[|f\hat{w}^{\tau} - fw| > \delta fw_{< \tau}] = O(P).$$

The above constants are not optimized, but with pure O -notation, it is not as easy to make a formally clear statement. Ignoring the constants and the restriction $\delta < 2$, we see that our error bound for priority sampling with threshold τ is of the same type as the one in (15) for threshold sampling with fixed threshold $t = \tau$. In our case, the priority threshold τ is variable, and from Lemma 5 it follows that our probabilistic error bound

$$6\delta(f\hat{w}_{< \tau}/(3\tau), P) fw_{< \tau}$$

decreases with the priority threshold τ .

The proof of Theorem 6 is rather convoluted. With some target error probability P , we will identify t_{\min}, t_{\max} such that

- (i) With probability $1 - O(P)$, the priority threshold $\tau \in [t_{\min}, t_{\max}]$.

- (ii) With probability $1 - O(P)$, with a single random choice of the h_i but simultaneously for all $t \in [t_{\min}, t_{\max}]$, if $\delta = 6 \delta(f\hat{w}_{<t}/(3t), P) < 2$, then $|f\hat{w}^\tau - fw| \leq \delta fw_{<t}$.

A union bound on $\neg(i) \vee \neg(ii)$ implies Theorem 6. What makes (ii) very tricky to prove is that $\delta(f\hat{w}_{<t}/(3t), P)$ can vary a lot for different $t \in [t_{\min}, t_{\max}]$.

Priority confidence intervals The format of Theorem 6 makes it very easy to derive confidence intervals like those for threshold sampling. A priority sample with priority threshold τ gives us the exact value $fw_{\geq\tau}$ for weights at least as big as τ , and an estimate $f\hat{w}_{<\tau}^\tau$ for those below. For an upper bound on $fw_{<\tau}$, we compute

$$f\hat{w}_{<\tau}^+ = \max\{x \mid \delta = 6 \delta(x/(3\tau), P) \wedge (1 - \delta)x \leq f\hat{w}_{\geq\tau}^\tau\}.$$

Note that here in the upper bound, we only consider $\delta \leq 1$, so we do not need to worry about the restriction $\delta < 2$ in Theorem 6. For the lower bound, we use

$$f\hat{w}_{<\tau}^- = \min\{x \mid \delta = 6 \delta(x/(3\tau), P) < 2 \wedge (1 + \delta)x \geq f\hat{w}_{\geq\tau}^\tau\}.$$

Here in the lower bound, the restriction $\delta = 6 \delta(x/(3\tau), P) < 2$ prevents us from deriving a lower bound $x = f\hat{w}_{<\tau}^- \leq f\hat{w}_{\geq\tau}^\tau/3$. In such cases, we use the trivial lower bound $x = f\hat{w}_{<\tau}^- = 0$ which in distance from $f\hat{w}_{<\tau}^\tau$ is at most 1.5 times bigger. Now, by Theorem 6,

$$\Pr[fw_{\geq\tau} + f\hat{w}_{<\tau}^- \leq fw \leq fw_{\geq\tau} + f\hat{w}_{<\tau}^+] \geq 1 - O(P).$$

In cases where the exact part $fw_{\geq\tau}$ of an estimate is small compared with the variable part $f\hat{w}_{\geq\tau}^\tau$, we may be interested in a non-zero lower bound $f\hat{w}_{<\tau}^-$ even if it is smaller than $f\hat{w}_{\geq\tau}^\tau/3$. To do this, we need bounds for larger δ .

Large errors We are now going to present bounds that works for arbitrarily large relative errors δ . We assume a basic error probability function \wp satisfying Definition 4 (a) and (b) while (c) may not be satisfied.

The bounds we get are not as clean as those from Theorem 6, but we include them to show that something can be done also for $\delta \geq 2$. In particular, this means that we only worry about positive errors.

Theorem 7 *For items $i \in I$, let be given a weight vector $(w_i)_{i \in I}$ with corresponding fractions $(f_i)_{i \in I}$, a target error probability P , and a sample size k . Based on $(w_i)_{i \in I}$, let t_{\max} the smallest upper bound on a random priority threshold that is exceeded with probability at most P , that is, the probability of generating at least $k + 1$ priorities above t_{\max} is at most P . Consider a random priority sampling event. Set*

$$\begin{aligned} \ell &= 1 + \log_2(t_{\max}/\tau) \\ \delta &= \delta(fw_{<\tau}/\tau, P/\ell^2). \end{aligned}$$

Then

$$\Pr[f\hat{w}_{<\tau}^\tau > (2 + 2\delta)fw_{<\tau}] = O(P)$$

Complementing Theorem 6, we only intend to use Theorem 7 for large errors where $(1 + 2\delta) = O(\delta)$. We wish to provide a probabilistic lower bound for $fw_{<\tau}$. Unfortunately, we do not know t_{\max} which depends on the whole weight vector $(w_i)_{i \in I}$. However, based our priority sample, it is not hard to generate

a probabilistic upper bound \bar{t}_{\max} on t_{\max} such that $\Pr[\bar{t}_{\max} < t_{\max}] \leq P$. We then compute $\bar{\ell} = 1 + \log_2(\bar{t}_{\max}/\tau)$ and set

$$f\hat{w}_{<\tau}^- = \min\{x \mid \bar{\delta} = \delta(x/\tau, P/\bar{\ell}^2) \wedge 2(1 + \bar{\delta})x \geq f\hat{w}_{<\tau}^\tau\}.$$

Then by Theorem 7,

$$\Pr[f w_{<\tau} \geq f\hat{w}_{<\tau}^-] = 1 - O(P).$$

To see this, let $f\hat{w}_{<\tau}^*$ be the value we would have obtained if we had computed $f\hat{w}_{<\tau}^-$ using the real t_{\max} . Our error event is that $\bar{t}_{\max} < t_{\max}$ or $f w_{<\tau} < f\hat{w}_{<\tau}^*$. The former happens with probability at most P , and Theorem 7 states that the latter happens with probability $O(P)$. Hence none of these error events happen with probability $1 - O(P)$, but then $\bar{t}_{\max} \geq t_{\max}$, implying $f\hat{w}_{<\tau}^- \leq f\hat{w}_{<\tau}^* \leq f w_{<\tau}$.

Theorem 6, our main result, is proved in Sections 3.3–3.5. Theorem 7, which is much easier, is proved in Section 3.6. In Section 3.7 we show how to compute the \bar{t}_{\max} used for confidence lower bound with Theorem 7.

3.3 The priority threshold

To prove Theorem 6 we need a handle on the variable priority threshold, identifying the previously mentioned probabilistic lower and upper bounds, t_{\min} and t_{\max} .

With priority sampling we specify the number k of samples, and use as threshold τ the $(k+1)$ th priority. Recall for any threshold t that the subscript $\leq t$ indicates restriction to items with weight below t . To relate this notation to a priority sample of some specified size k , we let $k_{\leq t}$ denote k minus the number of items with weight bigger than t . We define $k_{<t}$ accordingly. With threshold t , the expected number of samples is $E[|I^t|] = k - k_{\leq t} + w_{\leq t}/t = k - k_{<t} + w_{<t}/t$. The last equality is because weights $w_i = t$ cancel out.

The ideal threshold We define the *ideal threshold* t^* to be the one leading to an expected number of exactly k samples, that is $w_{\leq t^*} = t^* k_{\leq t^*}$.

Lemma 8 $tk_{\leq t} - w_{\leq t}$ is strictly increasing in t , so (a) t^* is uniquely defined, (b) $w_{\leq t} > k_{\leq t}t$ for any $t < t^*$, and (c) $w_{\leq t} < k_{\leq t}t$ for any $t > t^*$.

Proof If t increases without passing any weight value, then $k_{\leq t}$ and $w_{\leq t}$ are not changed, and the statement is trivial. When t reaches the value of some weight w_i , then both $tk_{\leq t}$ and $w_{\leq t}$ are increased by the same value w_i (if there are j weights with value w_i , the increase is by jw_i). ■

We would like to claim that the priority sampling threshold τ is concentrated around t^* , but this may be far from true. To illustrate what makes things tricky to analyze, consider the case where, say, we have $k-1$ weights of size t^* , and then a lot of small weights that sum to t^* . In this case we get an upper bound on τ which is close to t^* , but we do not get any good lower bound on τ even if we have full randomness. On the other hand, in this case, it is only little weight that is affected by the downwards variance in τ .

Threshold upper bound We now present a certificate that we will use to bound the probability that the priority threshold exceeds a given threshold T . The certificate uses an extra parameter $t \leq T$ which is a concrete weight value. In order to get a priority threshold as big as T , we need $|I^T| \geq k+1$, and this requires $|I_{\leq t}^T| \geq k_{\leq t} + 1$. The expected value of $|I_{\leq t}^T|$ is $\mu_{\leq t}^T = E[|I_{\leq t}^T|] = w_{\leq t}/T$. We define $\delta_{\leq t}^{\uparrow T}$ such that

$$(1 + \delta_{\leq t}^{\uparrow T})\mu_{\leq t}^T = k_{\leq t} + 1.$$

The error event is thus equivalent to $|I_{\leq t}^T| \geq (1 + \delta_{\leq t}^{\uparrow T})\mu_{\leq t}^T$. The probability of the upper bound error event is bounded by $\wp(\mu_{\leq t}^T, \delta_{\leq t}^{\uparrow T})$. We say that t and T *witness that T is a threshold upper bound with error probability P* if

$$\wp^{\uparrow}(t, T) = \wp(\mu_{\leq t}^T, \delta_{\leq t}^{\uparrow T}) \leq P. \quad (17)$$

It may seem wasteful that we make the conservative assumption that all weights between t and T get priorities at least T , but this is because we only want to bound probabilities using the simple function $\wp(\mu, \delta)$.

Threshold lower bound We also need to introduce a lower bound t_{\min} for the priority threshold such that the probability that it is violated is bounded by P . In our later analysis of error probabilities bounds, it will be crucial that we only need to consider thresholds $t \geq t_{\min}$.

The basic condition we use here is that the priority threshold τ is smaller than or equal to T if and only if $|I_{\leq T}^T| \leq k_{\leq T}^T$. The expected value of $I_{\leq T}^T$ is $\mu_{\leq T}^T = E[I_{\leq T}^T] = w_{\leq T}/T$. We define $\delta_{\leq T}^{\downarrow}$ such that

$$(1 - \delta_{\leq T}^{\downarrow})\mu_{\leq T}^T = k_{\leq T}^T.$$

The error event $\tau \leq T$ is thus equivalent that $|I_{\leq T}^T| \leq (1 - \delta_{\leq T}^{\downarrow})\mu_{\leq T}^T$. The probability of the lower bound error event is bounded by $\wp(\mu_{\leq T}^T, \delta_{\leq T}^{\downarrow})$. We say that T *witness that T is a threshold lower bound with error probability P* if

$$\wp^{\downarrow}(T) = \wp(\mu_{\leq T}^T, \delta_{\leq T}^{\downarrow}) \leq P. \quad (18)$$

Below t_{\min} is defined as the largest value witnessed this way. In the analysis below it will be crucial that t_{\max} is no bigger than any upper bound witness and that t_{\min} is no smaller than any lower bound witness.

By definition, $\Pr[\tau \in (t_{\min}, t_{\max})] \geq 1 - 2P$. To prove Theorem 6 it therefore suffices to show that the following good non-error event happens with probability $1 - O(P)$:

$$\begin{aligned} \forall t \in (t_{\min}, t_{\max}) : \delta = \delta(f\hat{w}_{<t}/(3t), P) &< 1/3 \\ \implies |f\hat{w}^t - fw| &> 6\delta fw_{<t} \end{aligned} \quad (19)$$

Note that our variable δ is 6 times smaller than the one in the statement of Theorem 6. This turns out to be a more convenient choice for the analysis.

3.4 Tightening the gap

The following lemmas give us a much tighter understanding of $t \in (t_{\min}, t_{\max})$.

Lemma 9 *If $t \geq t^*$ then $t(1 + \delta(w_{\leq t}/t_{\max}, P)) \geq t_{\max}$.*

Proof Let $\delta = \delta(w_{\leq t}/t_{\max}, P)$ and $T = (1 + \delta)t$. Assume for a contradiction that $T < t_{\max}$. We will prove that $\wp^{\uparrow}(t, T) < P$, hence that T is an upper bound, contradicting the minimality of t_{\max} .

By Lemma 8 (c), since $t \geq t^*$, we have $w_{\leq t} \leq tk_{\leq t}$, so

$$k_{\leq t} \geq w_{\leq t}/t = (1 + \delta)w_{\leq t}/T$$

For the priority threshold to be as big as T we need $|I_{\leq t}^T| \geq k_{\leq t} + 1$, and we defined $\delta_{\leq t}^{\uparrow T}$ such that $(1 + \delta_{\leq t}^{\uparrow T})w_{\leq t}/T = k_{\leq t} + 1$. It follows that $\delta_{\leq t}^{\uparrow T} > \delta$. Moreover, since $T < t_{\max}$, we have $w_{\leq t}/T > w_{\leq t}/t_{\max}$, so conclude

$$\wp^{\uparrow}(t, T) = \wp(w_{\leq t}/T, \delta_{\leq t}^{\uparrow T}) \leq \wp(w_{\leq t}/t_{\max}, \delta) < P.$$

It follows that T is an upper bound contradicting the minimality of t_{\max} . \blacksquare

Lemma 10 For $t \in [t_{\min}, t^*]$ then $t \geq (1 - \delta(w_{\leq t}/t, P))t^*$.

Proof The proof is by contradiction against our maximal lower bound t_{\min} . Recall that the priority threshold is bigger than T if $|I_{\leq T}^T| \leq k_{\leq T}$. The expectancy is $E[|I_{\leq T}^T|] = w_{\leq T}/T$, so for lower bounds, we defined $\delta_{\leq T}^\downarrow$ so that $(1 - \delta_{\leq T}^\downarrow)w_{\leq T}/T = k_{\leq T}$, getting T as a lower bound if $\wp^\downarrow(T) = \wp(w_{\leq T}/T, \delta_{\leq T}^\downarrow) \leq P$.

Suppose for a contradiction that $(1 - \delta)w_{\leq t}/t > k_{\leq t}$. Then, by definition, $\delta_\ell \leq \delta_{\leq t}^\downarrow$. It follows that

$$\wp^\downarrow(t) = \wp(w_{\leq t}/t, \delta_{\leq t}^\downarrow) < \wp(w_{\leq t}/t, \delta) \leq P,$$

implying that $t \geq t_{\min}$ is a lower bound. If $t > t_{\min}$ this contradicts the maximality of t_{\min} . Otherwise, we pick an infinitesimally larger $t' > t$ with no weights in $(t, t']$, that is, $w_{\leq t'} = w_{\leq t}$. This yields an infinitesimally larger $\wp^\downarrow(t')$, that is, $\wp^\downarrow(t) < \wp^\downarrow(t') < P$. Now $t' > t_{\min}$ is a lower bound contradicting the maximality of t_{\min} . Thus we conclude that $(1 - \delta)w_{\leq t}/t \leq k_{\leq t}$, or equivalently, $t > (1 - \delta)w_{\leq t}/k_{\leq t}$. Finally by Lemma 8 (b), we have $w_{\leq t}/k_{\leq t} \geq t^*$. This completes the proof that $t \geq (1 - \delta_\ell)t^*$. \blacksquare

In order to give a joint analysis for t bigger and smaller than t^* , we make a conservative combination of Lemma 9 and 10.

Lemma 11 Suppose $(t^-, t^+) = [t_{\min}, t^*]$ or $(t^-, t^+) = (t^*, t_{\max}]$. If $t \in (t^-, t^+]$ then $t \geq (1 - \delta(w_{\leq t}/t^+, P))t^+$.

Proof Let $\delta = \delta(w_{\leq t}/t^+, P)$. If $(t^-, t^+) = (t^*, t_{\max}]$, by Lemma 9, $t \geq t^+/(1 + \delta) \geq (1 - \delta)t^+$. If $(t^-, t^+) = [t_{\min}, t^*]$, by Lemma 10, $t \geq (1 - \delta(w_{\leq t}/t, P))t^+ > (1 - \delta)t^+$. The last inequality uses that $\delta(w_{\leq t}/t, P) \leq \delta$. This is because $w_{\leq t}/t \geq w_{\leq t}/t^+$, so by Definition 4 (a), $\wp(w_{\leq t}/t, \delta) \geq \wp(w_{\leq t}/t^+, \delta) = P$. \blacksquare

Loosing a factor 2 in the error probability to cover t bigger or smaller than t^* , our good event (19) reduces to

$$\begin{aligned} \forall t \in (t^-, t^+] : \delta = \delta(f\hat{w}_{< t}/(3t), P) &< 1/3 \\ \implies |f\hat{w}^t - fw| &> 6\delta fw_{< t} \end{aligned} \quad (20)$$

By Lemma 11, the bound on δ implies that $t \geq (2/3)t^+$. Therefore (20) is implied by

$$\begin{aligned} \forall t \in (t^-, t^+] : \delta = \delta(f\hat{w}_{< t}/(2t^+), P) &< 1/3 \\ \implies |f\hat{w}^t - fw| &\leq 6\delta fw_{< t} \end{aligned} \quad (21)$$

One advantage of dealing with $f\hat{w}_{< t}/t^+$ instead of $f\hat{w}_{< t}/t$ is that $f\hat{w}_{< t}/t^+$ is proportional to $f\hat{w}_{< t}$ hence monotone in t whereas $f\hat{w}_{< t}/t$ is not monotone in t . In particular, if there is a $t \in (t^-, t^+]$ such that $\delta(f\hat{w}_{< t}/(2t^+), P) \geq 1/3$, then there is a t'^- such that $\delta(f\hat{w}_{< t'}/(3t^+), P) < 1/3$ if and only if $t > t'^-$. Otherwise we set $t'^- = t^-$. Then (21) is equivalent to

$$\begin{aligned} \forall t \in (t'^-, t^+] : \delta = \delta(f\hat{w}_{< t}/(2t^+), P) \\ \implies |f\hat{w}^t - fw| &\leq 6\delta fw_{< t} \end{aligned} \quad (22)$$

and for $t \in (t'^-, t^+]$, we know that $\delta(f\hat{w}_{< t}/(2t^+), P) < 1/3$.

3.5 Dividing into layers

We now define a sequence $t_0 > t_1 > \dots > t_{L+1}$ of decreasing thresholds with $t_0 = t^+$ and $t_{L+1} = t'^-$. For $\ell = 0, \dots, L$ we require

$$fw_{\leq t_{\ell+1}} \geq fw_{< t_\ell}/2. \quad (23)$$

For $\ell < L$, we pick $t_{\ell+1}$ smallest possible satisfying (23), so

$$fw_{< t_{\ell+1}} < fw_{< t_\ell}/2, \quad (24)$$

We arrive $\ell = L$ when $fw_{\leq t'^-} \geq fw_{< t_L}/2$. Then we set $t_{L+1} = t'^-$.

For each “layer” $\ell \leq L$, we define

$$\delta_\ell = \delta(fw_{< t_\ell}/(2t^+), P), \quad (25)$$

noting that this is the same value as we would use for $t = t_\ell$ in (22). By definition, for all $\ell \leq L$, we have $\wp(fw_{< t_\ell}/(2t^+), \delta_\ell) = P$. Since $t_\ell > t'^-$, we have $\delta_\ell \leq 1/3 \leq 1$, so it follows from Definition 4 (c) that there is a constant C such that $fw_{< t_\ell} \delta_\ell^2 = C$ for all $\ell \leq L$. For $\ell = 1, \dots, L$, by (24), we have $fw_{< t_\ell} < fw_{< t_{\ell-1}}/2$. Therefore

$$\delta_\ell > \sqrt{2} \delta_{\ell-1} \quad (26)$$

$$\delta_\ell fw_{< t_\ell} < \delta_{\ell-1} fw_{< t_{\ell-1}}/\sqrt{2} \quad (27)$$

This will correspond to an effect where the relative errors are geometrically increasing while the absolute errors are geometrically decreasing. Another important thing to notice is that by (23), $fw_{\leq t_{\ell+1}} \geq fw_{< t_\ell}/2$, so $\delta(fw_{< t_{\ell+1}}/t^+, P) \leq \delta(fw_{< t_\ell}/(2t^+), P) = \delta_\ell$. Therefore, by Lemma 11,

$$t_{\ell+1} \geq (1 - \delta_\ell)t^+. \quad (28)$$

Good layers For each layer $\ell < L$, our good event will be that for weights in $[t_{\ell+1}, t_\ell)$, the relative estimate error is bounded by $2\delta_\ell$. Formally

$$\begin{aligned} \forall t \in (t_{\ell+1}, t^+] : & \left| f\hat{w}_{[t_{\ell+1}, t_\ell)}^t - fw_{[t_{\ell+1}, t_\ell)} \right| \\ & \leq 2\delta_\ell fw_{[t_{\ell+1}, t_\ell)} \end{aligned} \quad (29)$$

Above, the subscript $[t_{\ell+1}, t_\ell)$ denotes the restriction to items i with weights $w_i \in [t_{\ell+1}, t_\ell)$. The last layer L is special in that we want to consider all weights below t_L . Here the good event is that

$$\begin{aligned} \forall t \in (t_{L+1}, t^+] : & \left| f\hat{w}_{< t_L}^t - fw_{< t_L} \right| \\ & \leq 6\delta_L fw_{< t_L} \end{aligned} \quad (30)$$

To prove Theorem 6, we are going to prove two statements.

- Assume that all layers are good satisfying (29) and (30). If for any $t \in (t'^-, t^+]$ we add up the errors from all relevant layers, then the total error is bounded by $6\delta(f\hat{w}_{< t}/(2t^+), P)fw_{< t}$, so (22) satisfied.
- If P_ℓ is the probability that a layer ℓ fails, then the P_ℓ are geometrically increasing and $P_L = O(P)$, so by union, the probability that any layer fails is $O(P)$.

Adding layer errors Assuming that all layers are good satisfying (29) and (30), we pick an arbitrary threshold $t \in (t'^-, t^+]$. We wish to bound the estimate error $|f\hat{w}^t - fw|$.

Let h be the layer such that $t \in (t_{h-1}, t_h]$. We can only have estimate errors from weights below $t \leq t_h$, so

$$\begin{aligned}
|f\hat{w}^t - fw| &\leq \sum_{\ell=h}^{L-1} \left| f\hat{w}_{[t_{\ell+1}, t_{\ell})}^t - fw_{[t_{\ell+1}, t_{\ell})} \right| \\
&\quad + |f\hat{w}_{< t_L}^t - fw_{< t_L}| \\
&\leq \sum_{\ell=h}^{L-1} 2\delta_{\ell} w_{[t_{\ell+1}, t_{\ell})} \\
&\quad + 3\delta_L fw_{< t_L} \\
&= \sum_{\ell=h}^{L-1} 2\delta_{\ell} (fw_{< t_{\ell}} - fw_{< t_{\ell+1}}) \\
&\quad + 3\delta_L fw_{< t_L}
\end{aligned}$$

By (26), the δ_{ℓ} are increasing, so in the above sum, every $fw_{< t_{\ell}}$ appears with a positive coefficient. It follows that we could only get a larger sum if (24) was more than tight with $fw_{< t_{\ell}} = fw_{< t_{\ell-1}}/2$ for $\ell \leq L$. Then we would have $fw_{< t_{\ell}} - fw_{< t_{\ell+1}} = fw_{< t_{\ell}}/2$ and corresponding to (27), $\delta_{\ell} fw_{< t_{\ell}} = \delta_{\ell-1} fw_{< t_{\ell-1}}/\sqrt{2}$. Thus we get

$$\begin{aligned}
|f\hat{w}^t - fw| &\leq \sum_{\ell=h}^{L-1} 2\delta_{\ell} (fw_{< t_{\ell}} - fw_{< t_{\ell+1}}) \\
&\quad + 3\delta_L fw_{< t_L} \\
&\leq \sum_{\ell=h}^{L-1} \delta_h fw_{< t_h} / \sqrt{2}^{\ell-h} \\
&\quad + 3\delta_h fw_{< t_h} / \sqrt{2}^{L-h} \\
&< 3\delta_h fw_{< t_h}
\end{aligned}$$

The last inequality exploits that $\sum_{i=1}^{\infty} 1/\sqrt{2}^i = 1/(1 - 1/\sqrt{2}) < 3$. Since $t \leq t_h$, we have $\delta(fw_{< t}/(2t^+), P) \geq \delta(fw_{< t_h}/(2t^+), P) = \delta_h$. Also, $t > t_{h+1}$, so by (23), $fw_{< t} \geq fw_{\leq t_{h+1}} \geq fw_{< t_h}/2$, so

$$\delta(fw_{< t}/(2t^+), P) fw_{< t} \geq \delta_h fw_{< t_h}/2.$$

Therefore

$$|f\hat{w}^t - fw| \leq 3\delta_h fw_{< t_h} \leq 6\delta(fw_{< t}/(2t^+), P) fw_{< t}.$$

Thus we conclude that (22) follows from (29) and (30).

Intermediate layer error probabilities We now consider the intermediate layers $\ell = 0, \dots, L-1$. We want to show that the probability P_{ℓ} of violating (29) increases geometrically with ℓ , yet remains bounded by P . First we consider the upper bound part of (29)

$$\forall t \in (t_{\ell+1}, t^+] : f\hat{w}_{[t_{\ell+1}, t_{\ell})}^t \leq (1 + 2\delta_{\ell}) fw_{[t_{\ell+1}, t_{\ell})}. \quad (31)$$

We claim that it can never be violated. The worst that can happen is that every item i in the layer gets sampled, and the estimate is at most $f_i t^+$. However, the items all have weight at least $t_{\ell+1}$ and by (28), $t_{\ell+1} \geq (1 - \delta_\ell) t^+$. The increase is thus by at most a factor $t^+ / t_{\ell+1} \leq 1 / (1 - \delta_{\ell+1})$, which for $\delta_\ell < 1/3$ is less than $(1 + 2\delta_\ell)$. Thus (31) is satisfied regardless of the random choices.

We now consider the lower bound part of (29)

$$\forall t \in (t_{\ell+1}, t^+] : f \hat{w}_{[t_{\ell+1}, t_\ell]}^t \geq (1 - 2\delta_\ell) f w_{[t_{\ell+1}, t_\ell]}. \quad (32)$$

This event could happen. To bound the probability, we will focus on the loss $f w_{[t_{\ell+1}, t_\ell]} - f \hat{w}_{[t_{\ell+1}, t_\ell]}^t$. When bounding the loss, we do not consider the gain from possible over estimates of sampled items. We only consider the actual losses $f_i w_i$ from unsampled items i . In fact, we consider an item i lost if $q_i \leq t^+$. This includes any item unsampled with some threshold $t \leq t^+$. The loss for every threshold $t \leq t^+$ is thus bounded as

$$f w_{[t_{\ell+1}, t_\ell]} - f \hat{w}_{[t_{\ell+1}, t_\ell]}^t \leq \sum_{i: w_i \in (t_{\ell+1}, t_\ell]} [q_i \leq t^+] f_i w_i.$$

We know that $w_i \geq t_{\ell+1} \geq (1 - \delta) t^+$. Therefore

$$\begin{aligned} \Pr[q_i \leq t^+] &= \Pr[h_i \geq w_i / t^+] \leq \Pr[h_i \geq t_{\ell+1} / t^+] \\ &= 1 - t_{\ell+1} / t^+ = \delta_\ell. \end{aligned}$$

The expected loss from layer ℓ is thus bounded by

$$\sum_{i: w_i \in (t_{\ell+1}, t_\ell]} \delta_\ell f_i w_i = \delta_\ell f w_{[t_{\ell+1}, t_\ell]}.$$

For (32) to fail, we need a loss that is twice this big, that is, $2\delta_\ell f w_{[t_{\ell+1}, t_\ell]}$. We know that items i 's loss contribution $[q_i \leq t^+] f_i w_i$ depends only on h_i and that it is at most t^+ . The probability of violating (32) is therefore bounded by

$$\wp(\delta_\ell f w_{[t_{\ell+1}, t_\ell]} / t^+, 1) \leq \wp(\delta_\ell f w_{< t_\ell} / (2t^+), 1).$$

Let $\mu_\ell = \delta_\ell f w_{< t_\ell} / (2t^+)$. Then $P_\ell = \wp(\mu_\ell, 1)$ is our bound on the error probability. From (27), we know that $\delta_\ell f w_{< t_\ell} < \delta_{\ell-1} f w_{< t_{\ell-1}} / \sqrt{2}$, so $\mu_\ell < \mu_{\ell-1} / \sqrt{2}$. It follows from Definition 4 (c) that $\wp(\mu_\ell, 1) > \sqrt{2} \wp(\mu_{\ell-1}, 1)$, so the P_ℓ are geometrically increasing, and their sum is bounded by $3P_{L-1}$.

Finally, by definition (25), $\wp(f w_{< t_{L-1}} / (2t^+), \delta_{L-1}) = P$. To compare P with P_{L-1} , by Definition 4 (c), we compare $f w_{< t_{L-1}} / (2t^+) \delta_{L-1}^2$ with $\mu_{L-1} 1^2 = \delta_{L-1} f w_{< t_{L-1}} / (2t^+)$, and conclude that $P_{L-1} \leq \delta_{L-1} P \leq P/3$. The probability that any intermediate layer $\ell < L$ fails (32) is thus at most P . Since (32) was always satisfied, we conclude that (30) is satisfied for all layers $\ell < L$ with probability P .

The last layer We now consider items i with weights below t_L . On the upper bound side, the good event (30) states that

$$\forall t \in (t_{L+1}, t^+] : f \hat{w}_{< t_L}^t \leq (1 + 3\delta_L) f w_{< t_L}. \quad (33)$$

We will show that (33) fails with probability less than P . For an upper bound on the estimate with any threshold $t \in (t_{L+1}, t^+]$, we include item i if $q_i > t_{L+1}$, and if so, we give it at an estimate of $f_i t^+$ which is bigger than the sampled estimate with threshold $t \leq t^+$. The result is at most a factor t^+ / t_{L+1} bigger than in the sampled estimate with threshold t_{L+1} , and by (28), $t_{L+1} \geq (1 - \delta_L) t^+$. Thus, regardless of the random choices made, we conclude that

$$\forall t \in (t_{L+1}, t^+] : f \hat{w}_{< t_L}^t \leq f \hat{w}_{\leq t_L}^{t_{L+1}} / (1 - \delta_L).$$

Consider the following error event:

$$f\hat{w}_{\leq t_L}^{t_{L+1}} > (1 + \delta_L)fw_{\leq t_L} \quad (34)$$

The maximal item contribution to $f\hat{w}_{\leq t_L}^{t_{L+1}}$ is bounded by $t_L \leq t^+$, so the probability of (34) is bounded by

$$\wp(fw_{\leq t_L}/t^+, \delta_L) \leq \wp(fw_{\leq t_L}/(2t^+), \delta_L)/2 = P/2.$$

If (34) does not happen, then since $\delta_L \leq 1/3$,

$$\begin{aligned} \forall t \in (t_{L+1}, t^+] : f\hat{w}_{< t_L}^t &\leq (1 + \delta_L)fw_{\leq t_L}/(1 - \delta_L) \\ &\leq (1 + 3\delta_L)fw_{\leq t_L}, \end{aligned}$$

which is the statement of (33). We conclude that (33) fails with probability at most $P/2$.

We now consider the lower bound side of (30) which states that

$$\forall t \in (t_{L+1}, t^+] : f\hat{w}_{< t_L}^t \geq (1 - 3\delta_L)fw_{< t_L} \quad (35)$$

The analysis is very symmetric to the upper bound case. For a lower bound for weights $w_i < t_L$ with any threshold $t \in (t_{L+1}, t^+]$, we only include i if $q_i > t^+$, and if so, we only give i the estimate $f_i t_{L+1}$ which is smaller than the sampled estimate with threshold $t > t_{L+1}$. Our samples are exactly the same as those we would pick with threshold t^+ , and our estimates are smaller by a factor $t_{L+1}/t^+ \geq (1 - \delta_L)$, so we conclude that regardless of the random choices,

$$\forall t \in (t_{L+1}, t^+] : f\hat{w}_{< t_L}^t \geq f\hat{w}_{< t_L}^{t^+} (1 - \delta_L)$$

We now consider the following error event:

$$f\hat{w}_{\leq t_L}^{t^+} < (1 - \delta_L)fw_{\leq t_L} \quad (36)$$

The maximal item contribution to $f\hat{w}_{\leq t_L}^{t^+}$ is t^+ , so as for (34), we get that the probability of (36) is bounded by $\wp(fw_{\leq t_L}/t^+, \delta_L) \leq P/2$.

If (36) does not happen, then since $\delta_L \leq 1/3$,

$$\begin{aligned} \forall t \in (t_{L+1}, t^+] : f\hat{w}_{< t_L}^t &\geq (1 - \delta_L)fw_{\leq t_L}(1 - \delta_L) \\ &> (1 - 2\delta_L)fw_{\leq t_L}, \end{aligned}$$

which implies (35). We conclude that (33) fails with probability at most $P/2$. Including the probability of an upper bound error, we get that (30) fails with probability P .

Summing up Above we proved that the probability that (29) failed for any layer $\ell < L$ was at most P . We also saw that (30) failed with probability P . If none of them fail, we proved that (22) and hence (20) was satisfied, so (20) fails with probability at most $2P$. However, for (19) we need (20) both for $(t^-, t^+) = [t_{\min}, t^*]$ and for $(t^-, t^+) = (t^*, t_{\max}]$, so (19) fails with probability at most $4P$. Finally, we need to consider both the case that $\tau > t_{\max}$ and $\tau \leq t_{\min}$. Either of these events happens with probability at most P , so we conclude that the overall error probability is at most $6P = O(P)$. If no error happened and $\delta = \delta(f\hat{w}_{< t}/(3t), P) \leq 1/3$, then $|f\hat{w}^t - fw| > 6\delta fw_{< t}$. This completes the proof of Theorem 6.

3.6 Large errors

The limitation of Theorem 6 is that it can only be used to bound the probability that the estimate error $|f\hat{w}^\tau - f w_{<\tau}|$ is bigger than $2f w_{<\tau}$. Note here that errors above $f w_{<\tau}$ can only be overestimates. Since we are targeting arbitrarily large relative errors δ , for the probability function \wp , we can only assume conditions (a) and (b) in Definition 4.

We will use some of the same ideas as we used for the last layer in Section 3.5, but tuned for our situation. We will study intervals based on $t_\ell = t_{\max}/2^\ell$ for $\ell = 1, 2, \dots$. Interval ℓ is for thresholds $t \in (t_\ell, t_{\ell-1}] = (t_\ell, 2t_\ell]$. To define the error for interval ℓ , set

$$\delta_\ell = \wp(f w_{<t_\ell}/t_\ell, P/\ell^2).$$

The good non-error event for interval ℓ is that

$$f\hat{w}_{<t_\ell}^{t_\ell} \leq (1 + \delta_\ell) f w_{<t_\ell}. \quad (37)$$

By definition, the probability that (37) is violated is at most P/ℓ^2 , so the probability of failure for any ℓ is bounded by $\sum_{\ell=1}^{\infty} P/\ell^2 \leq P\pi^2/6 < 1.65P$. By definition, the probability that

$$\tau \leq t_{\max} \quad (38)$$

is violated is at most P , so our total error probability is bounded by $2.65P$. Below we assume no errors, that is, (37) holds for all ℓ and so does (38).

Consider an arbitrary threshold $t \in (0, t_{\max}]$ and let ℓ be such that $t \in (t_\ell, 2t_\ell]$. We can only have errors for weights $w_i < t$, so we want an upper bound on $f\hat{w}_{<t}^t$. The basic idea for an upper bound is to say that we sample all items with priority above t_ℓ , just as in the estimate $f\hat{w}_{<t_\ell}^{t_\ell}$, but instead of giving sampled item i estimate $\max\{w_i, < t_\ell\}$, it gets value $\max\{w_i, t\}$ which is at most $t/t_\ell \leq 2$ times bigger. Thus, regardless of the random choices,

$$f\hat{w}_{<t}^t < 2f\hat{w}_{<t_\ell}^{t_\ell}.$$

Assuming no error as in (37), we get

$$f\hat{w}_{<t}^{t_\ell} = f w_{[t_\ell, t]} + f\hat{w}_{<t_\ell}^{t_\ell} \leq f w_{<t} + \delta_\ell f w_{<t_\ell}.$$

Let

$$\delta = \delta(w_{<t}/t, P/\ell^2)$$

By Lemma 5, since $t > t_\ell$, we have $\delta_\ell f w_{<t_\ell} < \delta f w_{<t}$, so

$$f\hat{w}_{<t}^{t_\ell} \leq f w_{<t} + \delta_\ell f w_{<t_\ell} < f w_{<t} + \delta f w_{<t_\ell}.$$

We thus conclude

$$\forall t \in (0, t_{\max}] : f\hat{w}_{<t}^t < 2f\hat{w}_{<t_\ell}^{t_\ell} < 2(1 + \delta)f w_{<t}, \quad (39)$$

By (38) this bound also holds for $t = \tau$. Relating to the parameters in Theorem 7, note that for given $t \leq t_{\max}$, we picked $\ell = 1 + \lceil \log_2 t_{\max} \rceil \leq 1 + \log_2 t_{\max}$, so our relative error δ is no bigger than that in Theorem 7. This completes the proof of Theorem 7.

3.7 Upper bounding the upper bound

In Theorem 7 we defined t_{\max} to be smallest upper bound on a random priority threshold that is exceeded with probability at most P , that is, the probability of generating at least $k + 1$ priorities above t_{\max} is at most P . As described in Section 3.2, in applications of Theorem 7, it suffices if we based on our sample can compute a probabilistic upper bound \bar{t}_{\max} on the upper bound t_{\max} such that

$$\bar{t}_{\max} \geq t_{\max} \quad (40)$$

with probability at least $1 - P$. For better confidence lower bounds, we want \bar{t}_{\max} to be small. If the weight vector $(w_i)_{i \in I}$ was known, then we could compute $\bar{t}_{\max} \geq t_{\max}$ as described in Section 3.3.

When we do our priority sample and get priority threshold τ , we get all weights above τ exactly. The idea is now to compute a probabilistic upper bound $\bar{w}_{<\tau}$ on $w_{<\tau}$ such that

$$\bar{w}_{<\tau} \geq w_{<\tau}. \quad (41)$$

with probability at least $1 - P$. Given $\bar{w}_{<\tau}$, for any $t \geq \tau$, we define

$$\bar{w}_{<t} = w_{[\tau,t)} + \bar{w}_{<\tau}$$

If (41) is true, then $\bar{w}_{<t} \geq \bar{w}_{<t}$ for all $t \geq \tau$. Using these values in the upper bound from Section 3.3, we get a \bar{t}_{\max} satisfying (40).

It remains to derive a probabilistic upper bound $\bar{w}_{<\tau}$ on the original weight below the priority threshold τ . Based on the (partially unknown) original weight vector $(w_i)_{i \in I}$, for any threshold value t , we define δ_t to be the largest value such that $\wp(k_{<t}/(1 - \delta_t), \delta_t) \leq P$. Moreover, define $\bar{\mu}_{<t}^t = k_{<t}/(1 - \delta_t)$ and $\bar{w}_{<t} = t\bar{\mu}_{<t}^t$. Note that all these values are functions of t and $k_{<t}$ that can be derived from the priority sample if $t \geq \tau$. The following lemma states that $\bar{w}_{<\tau}$ thus defined does give us the desired probabilistic upper bound on $w_{<\tau}$

Lemma 12 *For the random priority threshold τ , the probability that $w_{<\tau} > \bar{w}_{<\tau}$ is at most P .*

Proof For any given set of input weights consider a threshold t such that $\bar{w}_{<t} \leq w_{<t}$. We claim that the random priority threshold τ is expected no smaller than t . Note that $\tau < t$ if and only if $|I_{<t}^t| \leq k_{<t}$. Since $\bar{w}_{<t} \leq w_{<t}$, we have $E[|I_{<t}^t|] = w_{<t}/t \geq \bar{\mu}_{<t}$. Moreover, $k_{<\tau} = (1 - \delta_\tau)\bar{\mu}_\tau$, so

$$\begin{aligned} \Pr[\tau \leq t] &= \Pr[|I_{<t}^t| \leq k_{<t}] = \Pr[|I_{<t}^t| \leq (1 - \delta_t)\bar{\mu}_t] \\ &\leq \wp(\bar{\mu}_{<t}, \delta_t) = \wp(k_{<t}/(1 - \delta_t), \delta_t) \leq P. \end{aligned}$$

Let t^+ be the maximal value such that $\bar{w}_{<t^+} \leq w_{<t^+}$. The probability that $\tau \leq t^+$ is at most P , and if $\tau > t^+$, then $\bar{w}_{<\tau} > w_{<\tau}$. If the maximal value t^+ does not exist, we define instead t^+ as the limit where $\bar{w}_{<t} \leq w_{<t}$ for $t < t^+$ while $\bar{w}_{<t^+} > w_{<t^+}$. The probability that $\tau < t^+$ is at most P , and if $\tau \geq t^+$, we again have $\bar{w}_{<\tau} > w_{<\tau}$. ■

3.8 Histogram similarity

As mentioned earlier, with weighted items, given the priority samples of sets A and B , we can easily estimate the weight of their intersection and union. As for bottom- k sample, the point is that we can construct the priority sample of their union. This priority sample involves the top $k + 1$ priorities from $A \cup B$. If one of these is in A , then it must also be among the top $k + 1$ priorities from A .

However, if we want to compare histograms, then it is natural to say that the same item i may have different weights in different sets. The item i has weight w_i^A in A and weight w_i^B in B . Let $w_i^{\max} = \max\{w_i^A, w_i^B\}$ and $w_i^{\min} = \min\{w_i^A, w_i^B\}$. The histogram similarity is $w^{\min}/w^{\max} = (\sum_i w_i^{\min})/(\sum_i w_i^{\max})$.

This would seem a perfect application of our fractional subsets with $w_i = w_i^{\max}$ and $f_i = w_i^{\min}/w_i^{\max}$. The issue is as follows. From our priority samples over the w_i^A and w_i^B we can easily get the priority sample for the $w_i = w_i^{\max}$. However, for the i sampled, we would typically not have a sample with w_i^{\min} , and then we cannot compute f_i .

Our solution is to keep the instances of an item i in A and B separate as twins i^A and i^B with priorities $q_i^A = w_i^A/h_i$ and $q_i^B = w_i^B/h_i$. For coordination, we still use the same hash value h_i to determine the priorities. If $w_i^A = w_i^B$, we get $q_i^A = q_i^B$, and then we break the tie in favor of i^A . The priority sample for $A \cup B$ consists of the top k split items, and the priority threshold τ is the $k + 1$ biggest among all priorities. Estimation is done as usual. The important point here is the interpretation of the results. If $w_i^A \geq w_i^B$, then the priority of i^A is higher than that of i^B . Thus, in our sample, when we see an item i^C , $C \in \{A, B\}$, we count it for the union \hat{w}^{\max} if it is not preceded by its twin; otherwise we count it for the intersection \hat{w}^{\min} .

The resulting estimators \hat{w}^{\min} and \hat{w}^{\max} will no longer be unbiased with truly random hashing. However, for our probabilistic error bounds with pseudo-random hashing, there is no asymptotic effect to our analysis. All the current analysis is using union bounds over threshold sampling events, using the fact that each hash value h_i contributes at most 1 to the number of items with priorities above a given threshold t . Now h_i affects 2 twins, but this is fine since all we need is that the contribution of each random variable is bounded by a constant.

References

- [1] R. Adler, R. Feldman, and M. Taqqu. *A Practical Guide to Heavy Tails*. Birkhauser, 1998.
- [2] Y. Bachrach, R. Herbrich, and E. Porat. Sketching algorithms for approximating rank correlations in collaborative filtering systems. In *Proc. 16th SPIRE*, pages 344–352, 2009.
- [3] Y. Bachrach and E. Porat. Fast pseudo-random fingerprints. *CoRR*, abs/1009.5791, 2010.
- [4] Y. Bachrach, E. Porat, and J. S. Rosenschein. Sketching techniques for collaborative filtering. In *Proc. 21st IJCAI*, pages 2016–2021, 2009.
- [5] Z. Bar-Yossef, T. S. Jayram, R. Kumar, D. Sivakumar, and L. Trevisan. Counting distinct elements in a data stream. In *International Workshop on Randomization and Approximation Techniques in Computer Science (RANDOM)*, pages 1–10, 2002.
- [6] A. Z. Broder. On the resemblance and containment of documents. In *Proc. Compression and Complexity of Sequences (SEQUENCES)*, pages 21–29, 1997.

- [7] A. Z. Broder. Identifying and filtering near-duplicate documents. In *Proc. 11th CPM*, pages 1–10, 2000.
- [8] A. Z. Broder, M. Charikar, A. M. Frieze, and M. Mitzenmacher. Min-wise independent permutations. *Journal of Computer and System Sciences*, 60(3):630–659, 2000. See also STOC’98.
- [9] A. Z. Broder, S. C. Glassman, M. S. Manasse, and G. Zweig. Syntactic clustering of the web. *Computer Networks*, 29:1157–1166, 1997.
- [10] E. Cohen, M. Datar, S. Fujiwara, A. Gionis, P. Indyk, R. Motwani, J. D. Ullman, and C. Yang. Finding interesting associations without support pruning. *IEEE Trans. Knowl. Data Eng.*, 13(1):64–78, 2001.
- [11] E. Cohen, N. Duffield, H. Kaplan, C. Lund, and M. Thorup. Efficient stream sampling for variance-optimal estimation of subset sums. *SIAM Journal on Computing*, 40(5):1402–1431, 2011. Announced at SODA’09.
- [12] E. Cohen and H. Kaplan. Summarizing data using bottom-k sketches. In *Proc. 26th PODC*, pages 225–234, 2007.
- [13] M. Datar and S. Muthukrishnan. Estimating rarity and similarity over data stream windows. In *Proc. 10th ESA*, pages 323–334, 2002.
- [14] M. Dietzfelbinger. Universal hashing and k-wise independent random variables via integer arithmetic without primes. In *Proc. 13th STACS*, pages 569–580, 1996.
- [15] M. Dietzfelbinger, J. Gil, Y. Matias, and N. Pippenger. Polynomial hash functions are reliable (extended abstract). In *Proc. 19th ICALP*, pages 235–246, 1992.
- [16] N. Duffield, C. Lund, and M. Thorup. Learn more, sample less: control of volume and variance in network measurements. *IEEE Transactions on Information Theory*, 51(5):1756–1775, 2005.
- [17] N. Duffield, C. Lund, and M. Thorup. Priority sampling for estimation of arbitrary subset sums. *J. ACM*, 54(6):Article 32, 2007. Announced at SIGMETRICS’04.
- [18] G. Feigenblat, E. Porat, and A. Shifan. Exponential space improvement for minwise based algorithms. In *Proc. FSTTCS*, pages 70–85, 2012.
- [19] J. M. Hellerstein, P. J. Haas, and H. J. Wang. Online aggregation. In *Proc. SIGMOD*, pages 171–182, 1997.
- [20] M. R. Henzinger. Finding near-duplicate web pages: a large-scale evaluation of algorithms. In *Proc. 29th SIGIR*, pages 284–291, 2006.
- [21] C. P. Kruskal, L. Rudolph, and M. Snir. A complexity theory of efficient parallel algorithms. *Theor. Comput. Sci.*, 71(1):95–132, 1990.
- [22] G. S. Manku, A. Jain, and A. D. Sarma. Detecting near-duplicates for web crawling. In *Proc. 16th WWW*, pages 141–150, 2007.
- [23] M. Mitzenmacher and S. P. Vadhan. Why simple hash functions work: exploiting the entropy in a data stream. In *Proc. 19th ACM/SIAM Symposium on Discrete Algorithms (SODA)*, pages 746–755, 2008.

- [24] D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, and N. Weaver. Inside the slammer worm. *IEEE Security and Privacy Magazine*, 1(4):33–39, 2003.
- [25] R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge University Press, 1995.
- [26] A. Pagh, R. Pagh, and M. Ružić. Linear probing with constant independence. *SIAM Journal on Computing*, 39(3):1107–1120, 2009. See also STOC’07.
- [27] M. Pătraşcu and M. Thorup. On the k -independence required by linear probing and minwise independence. In *Proc. 36th ICALP, Part I, LNCS 6198*, pages 715–726, 2010.
- [28] M. Pătraşcu and M. Thorup. Twisted tabulation hashing. In *Proc. 23rd SODA*, pages 209–228, 2013.
- [29] C.-E. Särndal, B. Swensson, and J. Wretman. *Model Assisted Survey Sampling*. Springer, 1992.
- [30] S. Schleimer, D. S. Wilkerson, and A. Aiken. Winnowing: Local algorithms for document fingerprinting. In *Proc. SIGMOD*, pages 76–85, 2003.
- [31] J. P. Schmidt, A. Siegel, and A. Srinivasan. Chernoff-Hoeffding bounds for applications with limited independence. *SIAM Journal on Discrete Mathematics*, 8(2):223–250, 1995. See also SODA’93.
- [32] M. Szegedy. The DLT priority sampling is essentially optimal. In *Proc. 38th STOC*, pages 150–158, 2006.
- [33] M. Thorup. Confidence intervals for priority sampling. In *Proc. SIGMETRICS*, pages 252–263, 2006.
- [34] M. Thorup. Bottom- k and priority sampling, set similarity and subset sums with minimal independence. *CoRR*, 2013.
- [35] M. Thorup and Y. Zhang. Tabulation-based 5-independent hashing with applications to linear probing and second moment estimation. *SIAM Journal on Computing*, 41(2):293–331, 2012. Announced at SODA’04 and ALENEX’10.
- [36] M. N. Wegman and L. Carter. New classes and applications of hash functions. *Journal of Computer and System Sciences*, 22(3):265–279, 1981. See also FOCS’79.
- [37] H. Yang and J. P. Callan. Near-duplicate detection by instance-level constrained clustering. In *Proc. 29th SIGIR*, pages 421–428, 2006.